ABSTRACT

File availability in distributed file storage systems is improved by initially placing replicas of a file or other object on different ones of multiple devices using a first process. Subsequently, the placement of the replicas is improved by evaluating whether any replicas of a first file can be swapped with any replicas of a second file without a reduction in the combined file availability of the first and second files, and swapping a replica of the first file with a replica of the second file if the swapping results in no reduction in the combined file availability of the first and second files.